Tit-For-Tat Equilibria in Discounted Repeated Games with Private Monitoring

Hitoshi Matsushima[1]

Department of Economics, University of Tokyo[2]

April 24, 2007

Abstract

We investigate infinitely repeated games with imperfect private monitoring. We focus on a class of games where the payoff functions are additively separable and the signal for monitoring a player's action does not depend on the other player's action. Tit-for-tat strategies function very well in this class, according to which each player's action in each period depends only on the signal for the opponent's action one period before. With almost perfect monitoring, we show that even if the discount factors are fixed low, efficiency is approximated by a tit-for-tat Nash equilibrium payoff vector.

JEL Classification Numbers: C72, C73, D82, H41

Keywords: Infinitely Repeated Games, Private Monitoring, Tit-For-Tat Strategies, Fixed Discount Factors, Approximate Efficiency

[1] This research was supported by a Grant-In-Aid for Scientific Research (KAKENHI 15330036) from the Japan Society for the Promotion of Science (JSPS) and the Ministry of Education, Culture, Sports, Science and Technology (MEXT) of the Japanese government as well as a grant from the Center for Advanced Research in Finance (CARF) at the University of Tokyo.

[2] Department of Economics, University of Tokyo, Hongo, Bunkyo-ku, Tokyo 113-0033, Japan. E-mail: hitoshi@e.u-tokyo.ac.jp

1. Introduction

We investigate two-player infinitely repeated games with discounting. We assume that monitoring is *imperfect* in that each player cannot observe the opponent's action choice but can imperfectly monitor it by observing a noisy signal. This signal is randomly determined according to the probability function conditional on the opponent's action choice. We also assume that monitoring is *private* in that the signals that each player observes cannot be observed by the opponent.

We focus on a class of games where the payoff functions are *additively separable* and the signals for monitoring any player's action choices are *independent* of the other player's action choices. With this class restriction, we investigate the possibility that efficiency is approximated by a Nash equilibrium payoff vector when monitoring is *almost* perfect, and we arrive at the following affirmative answer. In the case of almost perfect private monitoring, there exists a simple *tit-for-tat* Nash equilibrium that approximately induces efficiency, according to which, each player's action choice in each period depends only on the signal for the opponent's action choice one period before. The main contribution of this paper is that this positive result holds *even if the discount factor is fixed and not very high*.

It is well known that if each player *perfectly* monitors the opponent's action choices and the discount factor is not very low, efficiency is achieved by a Nash equilibrium. See Fudenberg and Tirole (1991) and Osborne and Rubinstein (1994). In order to clarify the robustness with respect to monitoring ability, it is important to answer the question of whether efficiency is approximated by a Nash equilibrium payoff vector if monitoring is almost perfect but imperfect. In the case of *public* monitoring where the signals are observable to both players, it is now not difficult to answer in the affirmative due to previous works such as Green and Porter (1984); Abreu, Pearce, and Stachetti (1990), and Fudenberg, Levine, and Maskin (1986). In the case of private monitoring, however, this robustness issue remains unresolved and is much more substantial.

In fact, the previous works have given only partial answers for the private monitoring case. First, Sekiguchi (1997) showed an example in which efficiency can be approximated by a Nash equilibrium payoff vector when monitoring is almost perfect. Sekiguchi demonstrated an idea of construction with public randomization, which allowed equilibrium strategies to depend on histories in a non-recursive manner. Ely and Valimaki (2002) and Piccione (2002) demonstrated another idea of construction on the basis of recursive Markovian techniques, which was applicable to a class of games wider than that presented in Sekiguchi. By constructing Markovian strategies to which both the cooperative and non-cooperative actions are the best responses at all times, they showed that efficiency is approximated by a Nash equilibrium payoff vector when the discount factor is very close to unity. However, this result is not satisfactory because their proofs crucially depend on the assumption of almost *no* discounting. Hence, the robustness is regarded as an open question in the case when the discount factor is fixed and not very high. This is the main theme of this paper.

With our class restriction, by using only tit-for-tat strategies instead of the more complicated Markovian ones, we can demonstrate the same result as that shown by Ely, Valimaki, and Piccione. Tit-for-tat Nash equilibria have a useful property that the least upper bound of tit-for-tat Nash equilibrium payoffs for each player is *independent* of the discount factor. This implies that whenever the approximate efficiency holds with almost no discounting, then this holds even if the discount factor is not very high. This is precisely what we will show as the main theorem.

This paper is organized as follows. Section 2 describes the model. Section 3 defines tit-for-tat strategies. In Section 4, we characterize the tit-for-tat Nash equilibria and the least upper bounds. Section 5 presents the main theorem. Finally, Section 6 considers an example where we argue the degree of implicit collusion sustained *in the long run* by introducing the concept of *stationary distribution* of action profiles.

2. The Model

A *two-person component game* is defined by $(A_i, u_i)_{i\in\{1,2\}}$, where A_i denotes the finite set of actions for each player $i \in \{1,2\}$, $a_i \in A_i$, $A \equiv A_1 \times A_2$, $a \equiv (a_1, a_2) \in A$, $u_i : A \to R$, and $u_i(a)$ is the payoff for player i when the players choose the action profile $a \in A$. Let $\alpha_i : A_i \to [0,1]$ denote a *mixed* action for player i. Let Δ_i denote the set of mixed actions for player i.

Two noisy *signals* $\omega_1 \in \Omega_1$ and $\omega_2 \in \Omega_2$ occur after the players' action choices, where Ω_i denotes the finite set of possible ω_i, $\omega = (\omega_1, \omega_2)$, and $\Omega = \Omega_1 \times \Omega_2$. A signal profile $\omega \in \Omega$ is randomly determined according to the probability function $f(\cdot | a) : \Omega \to R_+$ conditional on $a \in A$. Let $f_i(\omega_i | a) \equiv \sum_{\omega_j \in \Omega_j} f(\omega | a)$, where $j \neq i$.

We assume that the payoff functions are *additively separable,* i.e.,

$$u_i(a) = v_i(a_i) + w_i(a_j) \quad \text{for all } i \in \{1,2\} \text{ and all } a \in A,$$

where $v_i : A_i \to R$ and $w_i : A_j \to R$. We assume that $f_i(\omega_i | a)$ does *not* depend on a_j. We write $f_i(\omega_i | a_i)$ instead of $f_i(\omega_i | a)$ and call $\omega_i \in \Omega_i$ *the signal for player i*'s *action choice*. An economic situation relevant to our model is the *voluntary contribution of public goods*. The players supply their public goods that are perfectly differentiated each other. An action $a_i \in A_i$ for each player i implies the amount of public good that player i produces. The production cost for player i's public good is given by $-v_i(a_i)$. Player i's benefit from opponent j's public good is given by $w_i(a_j)$.

Fix A_1, A_2, Ω_1, and Ω_2 arbitrarily. We define an *infinitely repeated game* by $\Gamma \equiv (u_i, f_i, \delta_i)_{i\in\{1,2\}}$, where $\delta_i \in (0,1)$ denotes the *discount factor* for player i. We allow the players to have *different* discount factors. Let $h(t) = (a(\tau), \omega(\tau))_{\tau=1}^{t}$ denote a *history* up to period t, where $a(\tau) = (a_1(\tau), a_2(\tau)) \in A$, and let $\omega(\tau) = (\omega_1(\tau), \omega_2(\tau)) \in \Omega$ denote the action profile and the signal profile in period t, respectively. Let $H = \{h(t) \,|\, t = 0,1,...\}$ denote the set of histories, where $h(0)$ is the null history. A *strategy* for player $i \in \{1,2\}$ is defined as $\sigma_i : H \to \Delta_i$, where $\sigma_i(h(t-1))$ is the mixed action for player i in period t when $h(t-1)$ occurs. Let $\sigma = (\sigma_1, \sigma_2)$ denote a strategy profile. The *payoff* for player i induced by σ in Γ is given by $v_i(\sigma) \equiv (1-\delta_i)E[\sum_{\tau=1}^{t} \delta_i^{\tau-1} u_i(a(\tau)) \,|\, \sigma, \Gamma]$, where $E[\cdot \,|\, \sigma, \Gamma]$ denotes the expectation when the players conform to σ in Γ. Let Σ_i denote the set of strategies for player i. A strategy profile σ is said to be a *Nash equilibrium* in Γ if

$$v_i(\sigma) \geq v_i(\sigma_i', \sigma_j) \text{ for all } i \in \{1,2\} \text{ and all } \sigma_i' \in \Sigma_i.$$

If monitoring is *imperfect* and *private*, i.e., each player observes *only* the signal for the opponent’s action choice, it is appropriate to replace the set of strategies Σ_i with a subset $\hat{\Sigma}_i \subset \Sigma_i$ that is defined as the set of strategies σ_i such that $\sigma_i(h(t))$ is independent of $(a_j(\tau), \omega_i(\tau))$ for all $\tau \in \{1,...,t\}$. Let $\hat{\Sigma} = \hat{\Sigma}_1 \times \hat{\Sigma}_2$. A strategy profile $\sigma \in \hat{\Sigma}$ is said to be a *Nash equilibrium in* Γ *with private monitoring* if $v_i(\sigma) \geq v_i(\sigma_i', \sigma_j)$ for all $i \in \{1,2\}$ and all $\sigma_i' \in \hat{\Sigma}_i$. Clearly, *a strategy profile is a Nash equilibrium in* Γ *with private monitoring if it is a Nash equilibrium in* Γ.

3. Tit-For-Tat Equilibria

For each $i \in \{1,2\}$, we arbitrarily fix two actions $a_i^* \in A_i$ and $a_i^{**} \in A_i$ for player i. Let $a^* = (a_1^*, a_2^*) \in A$ and $a^{**} = (a_1^{**}, a_2^{**}) \in A$. We assume that the action profile a^* is *efficient* in that $u_1(a^*) + u_2(a^*) \geq u_1(a) + u_2(a)$ for all $a \in A$. Further, we assume that for every $i \in \{1,2\}$, $v_i(a_i^*) < v_i(a_i^{**})$ and $w_i(a_j^*) > w_i(a_j^{**})$, i.e.,

(1) $$u_i(a_i^*, a_j) < u_i(a_i^{**}, a_j) \text{ and } u_i(a_i, a_j^*) > u_i(a_i, a_j^{**}) \text{ for all } a \in A.$$

In the voluntary contributions of public goods, this assumption implies that player i is producing a_i^* at a higher cost than a_i^{**} but in a manner more beneficial to opponent j.

We consider strategies that are *tit-for-tat* in that

(i) each player i only chooses a_i^* and a_i^{**} at all times, and

(ii) each player i's mixed action in any period $t \geq 2$ depends *only* on the signal

$$\omega_j(t-1) \text{ for opponent } j\text{'s action choice in the previous period } t-1.$$

Formally, a strategy σ_i for each player $i \in \{1,2\}$ is said to be *tit-for-tat* if for every $t \geq 1$ and every $h(t-1) \in H$,

$$\sigma_i(h(t-1))(a_i) = 0 \quad \text{for all} \quad a_i \in A_i \setminus \{a_i^*, a_i^{**}\}$$

and for every $t \geq 2$, every $h(t-1) \in H$, and every $h'(t-1) \in H \setminus \{h(t-1)\}$,

$$\sigma_i(h(t-1)) = \sigma_i(h'(t-1)) \quad \text{whenever} \quad \omega_j(t-1) = \omega'_j(t-1).$$

A tit-for-tat strategy σ_i is represented by (q_i, s_i) where $q_i \in [0,1]$, $s_i : \Omega_j \to [0,1]$,

$$\sigma_i(h(0))(a_i^*) = q_i, \quad \sigma_i(h(0))(a_i^{**}) = 1 - q_i,$$

for every $t \geq 2$ and every $h(t-1) \in H$,

$$\sigma_i(h(t))(a_i^*) = 1 - s_i(\omega_j(t)), \text{ and } \sigma_i(h(t))(a_i^{**}) = s_i(\omega_j(t)).$$

According to (q_i, s_i), player i chooses action a_i^* (action a_i^{**}) with probability q_i (probability $1-q_i$) in period 1 and action a_i^* (action a_i^{**}) with probability $s_i(\omega_j)$ (probability $1 - s_i(\omega_j)$) if he/she observes ω_j one period before. Note that any tit-for-tat strategy for player i belongs to the subset $\hat{\Sigma}_i$.

Let (q,s) denote a tit-for-tat strategy profile, where $q = (q_1, q_2)$ and $s = (s_1, s_2)$. We confine our attention to the tit-for-tat Nash equilibria (q,s) in Γ such that each player i chooses a_i^* with *positive* probabilities in every period $t \geq 2$, i.e., $s_i(\omega_j(t)) > 0$ for some $\omega_j \in \Omega_j$. Note that the necessary and sufficient condition for a tit-for-tat strategy profile to be a Nash equilibrium in Γ is that each player has no incentive to choose any other tit-for-tat strategy. Hence, it follows that *a tit-for-tat strategy profile is a Nash equilibrium in* Γ *with private monitoring if and only if it is a Nash equilibrium in* Γ. The set of tit-for-tat Nash equilibria also remains unchanged when monitoring is imperfect and public, that is, when the signals are observable to both players.

4. Characterization

The definition of tit-for-tat strategies, along with the assumptions on the payoff functions and the signals in Section 3, implies that in every period t, the incentive constraint of Nash equilibrium for each player i is *irrelevant* to the history other than $\omega_j(t-1)$; further, each player $i's$ current action $a_i(t)$ influences opponent $j's$ mixed action $\sigma_j(h(t))$ in the next period $t+1$ through the determination of $\omega_i(t)$. Since each player $i's$ choice of a_i^* is at all times one of the best responses to any tit-for-tat Nash equilibrium (q,s), it follows that for every $i\in\{1,2\}$,

(2)
$$v_i(q,s) = v_i(a_i) + (1-\delta)\{q_j w_i(a_j^*) + (1-q_j) w_i(a_j^{**})\}$$
$$+\sum_{\omega_i\in\Omega_i} \delta_i [w_i(a_j^*) s_j(\omega_i) + w_i(a_j^{**})\{1-s_j(\omega_i)\}] f_i(\omega_i \mid a_i),$$

and that the incentive constraint is replaced with the maximization of

(3) $$v_i(a_i) + \sum_{\omega_i \in \Omega_i} \delta_i [w_i(a_j^*) s_j(\omega_i) + w_i(a_j^{**})\{1 - s_j(\omega_i)\}] f_i(\omega_i \mid a_i)$$

with respect to $a_i \in A_i$. The definition of tit-for-tat Nash equilibrium *automatically* implies that both a_i^* and a_i^{**} maximize this value. Hence, we have proved the following proposition.

Proposition 1: *A tit-for-tat strategy profile* (q,s) *is a Nash equilibrium in* Γ *if and only if for every* $i \in \{1,2\}$,

(4) $$v_i(a_i^*) + \sum_{\omega_i \in \Omega_i} \delta_i [w_i(a_j^*) s_j(\omega_i) + w_i(a_j^{**})\{1 - s_j(\omega_i)\}] f_i(\omega_i \mid a_i^*)$$
$$= v_i(a_i^{**}) + \sum_{\omega_i \in \Omega_i} \delta_i [w_i(a_j^*) s_j(\omega_i) + w_i(a_j^{**})\{1 - s_j(\omega_i)\}] f_i(\omega_i \mid a_i^{**}),$$

and for every $a_i \in A_i$,

$$v_i(a_i^*) + \sum_{\omega_i \in \Omega_i} \delta_i [w_i(a_j^*) s_j(\omega_i) + w_i(a_j^{**})\{1 - s_j(\omega_i)\}] f_i(\omega_i \mid a_i^*)$$
$$\geq v_i(a_i) + \sum_{\omega_i \in \Omega_i} \delta_i [w_i(a_j^*) s_j(\omega_i) + w_i(a_j^{**})\{1 - s_j(\omega_i)\}] f_i(\omega_i \mid a_i).$$

From Proposition 1, it follows that there exists a tit-for-tat Nash equilibrium in Γ if and only if for every $i \in \{1,2\}$, there exists a function $\mu_i : \Omega_i \to R_+ \cup \{0\}$ such that

(5) $$v_i(a_i^*) - \sum_{\omega_i \in \Omega_i} \mu_i(\omega_i) f_i(\omega_i \mid a_i^*) = v_i(a_i^{**}) - \sum_{\omega_i \in \Omega_i} \mu_i(\omega_i) f_i(\omega_i \mid a_i^{**}),$$

(6) $$v_i(a_i^*) - \sum_{\omega_i \in \Omega_i} \mu_i(\omega_i) f_i(\omega_i \mid a_i^*) \geq v_i(a_i) - \sum_{\omega_i \in \Omega_i} \mu_i(\omega_i) f_i(\omega_i \mid a_i) \text{ for all}$$

$a_i \in A_i$, and

(7) $$0 \le \mu_i(\omega_i) \le \delta_i\{w_i(a_j^*) - w_i(a_j^{**})\} \text{ for all } \omega_i \in \Omega_i .$$

From the compactness and non-emptiness of the set of possible μ_i satisfying (5), (6), and (7), we can define a value $R_i(\delta_i) = R_i(\delta_i;\Gamma) \in R$ as

$$\max_{\mu_i:\Omega_i \to R_+ \cup \{0\}} \{u_i(a^*) - \sum_{\omega_i \in \Omega_i} \mu_i(\omega_i) f_i(\omega_i \mid a_i^*)\} \text{ subject to (5), (6), and (7).}$$

The following proposition shows that $R_i(\delta_i)$ is regarded as the least upper bound of the tit-for-tat Nash equilibrium payoffs for each player $i \in \{1,2\}$.

Proposition 2: *Suppose that there exists a tit-for-tat Nash equilibrium in* Γ*, i.e., for each* $i \in \{1,2\}$*, there exists a function* μ_i *that satisfies (5), (6), and (7). Then, there exists a tit-for-tat Nash equilibrium* $(\hat{q},\hat{s})$ *in* Γ *such that*

$$v_i(\hat{q},\hat{s}) = R_i(\delta_i) \text{ for all } i \in \{1,2\} .$$

Moreover, for every tit-for-tat Nash equilibrium (q,s) *in* Γ*,*

$$v_i(q,s) \le R_i(\delta_i) \text{ for each } i \in \{1,2\} .$$

Proof: From (2), for every tit-for-tat Nash equilibrium (q,s) and every $i \in \{1,2\}$,

$$v_i(q,s) \le u_i(a^*) - \sum_{\omega_i \in \Omega_i} \delta_i\{w_i(a_j^*) - w_i(a_j^{**})\}\{1 - s_j(\omega_i)\} f_i(\omega_i \mid a_i^*) ;$$

note that in the above equation, the function μ_i defined by

$$\mu_i(\omega_i) = \delta_i\{w_i(a_j^*) - w_i(a_j^{**})\}\{1 - s_j(\omega_i)\} \text{ for all } \omega_i \in \Omega_i$$

satisfies (5), (6), and (7). This implies that $v_i(q,s) \le R_i(\delta_i)$.

For each $i \in \{1,2\}$, there exists $\hat{\mu}_i$ satisfying (5), (6), and (7) such that

$$u_i(a^*) - \sum_{\omega_i \in \Omega_i} \hat{\mu}_i(\omega_i) f_i(\omega_i \mid a_i^*) = R_i(\delta_i).$$

We can construct a tit-for-tat Nash equilibrium $(\hat{q}, \hat{s})$ in a manner that for every $i \in \{1,2\}$,

$$\hat{q}_i = 1, \text{ and } \hat{s}_j(\omega_i) = 1 - \frac{\hat{\mu}_i(\omega_i)}{\delta_i\{w_i(a_j^*) - w_i(a_j^{**})\}} \text{ for all } \omega_i \in \Omega_i.$$

Clearly, $v_i(\hat{q}, \hat{s}) = R_i(\delta_i)$ for each $i \in \{1,2\}$. **Q.E.D.**

Remark 1 (Independence of Discount Factors): Given that the discount factors are sufficiently large, we can verify that the least upper bound $R_i(\delta_i;\Gamma)$ for each player i does not depend on δ_i as follows. We define $\tilde{R}_i \in R$ by

$$\max_{\mu_i:\Omega_i \to R_+ \cup \{0\}} \{u_i(a^*) - \sum_{\omega_i \in \Omega_i} \mu_i(\omega_i) f_i(\omega_i \mid a_i^*)\} \text{ subject to (5) and (6),}$$

where we must note that $\tilde{R}_i$ is independent of δ_i, and that $\tilde{R}_i \ge R_i(\delta_i)$. If $\delta_i\{w_i(a_j^*) - w_i(a_j^{**})\}$ is large enough for the restriction (7) not to be binding, then it holds that $\tilde{R}_i = R_i(\delta_i)$. Let $\tilde{\delta}_i$ denote the minimal discount factor δ_i such that $\tilde{R}_i = R_i(\delta_i)$. Since $R_i(\delta_i)$ is nondecreasing with respect to δ_i, we have shown that

$$R_i(\delta_i) = R_i(\tilde{\delta}_i) = \tilde{R}_i \text{ for all } \delta_i \ge \tilde{\delta}_i.$$

Remark 2 (Exchangeability): Proposition 1 implies that if (q,s) is a Nash equilibrium in a repeated game $\Gamma = (u_i, f_i, \delta_i)_{i\in\{1,2\}}$, then all tit-for-tat strategies for each player i are the best responses to (q_j, s_j) in any repeated game such that player i's discount factor is the same as that of Γ, i.e., δ_i. This implies that the tit-for-tat Nash equilibrium notion satisfies the following *strong* property of *exchangeability* across different games. Consider three repeated games given by $\Gamma \equiv (u_i, f_i, \delta_i)_{i\in\{1,2\}}$, $\Gamma' \equiv (u_i', f_i', \delta_i')_{i\in\{1,2\}}$, and $\Gamma'' = (u_1, f_1, \delta_1, u_2', f_2', \delta_2')$. *If* (q,s) *is a Nash equilibrium in* Γ *and* (q',s') *is a Nash equilibrium in* Γ'*, then* (q_1, s_1, q_2', s_2') *is a Nash equilibrium in* Γ''*, where the payoffs are unchanged in that*

$$v_1(q_1, s_1, q_2', s_2'; \Gamma'') = v_1(q', s'; \Gamma') \quad \textit{and} \quad v_2(q_1, s_1, q_2', s_2'; \Gamma'') = v_1(q, s; \Gamma).$$

5. Approximate Efficiency

We assume that for each $i \in \{1,2\}$,

$$v_i(a_i) - v_i(a_i^*) < \delta_i \{w_i(a_j^*) - w_i(a_j^{**})\} \text{ for all } a_i \in A_i \setminus \{a_i^*\}, \tag{8}$$

which is a necessary condition for (7), i.e., the existence of tit-for-tat Nash equilibria. Note that (8) is sufficient and (almost) necessary for the existence of the efficient tit-for-tat Nash equilibrium given that monitoring is *perfect*. We will show that (8) is also sufficient for the existence of the approximately efficient tit-for-tat Nash equilibrium given that monitoring is almost perfect.

Fix $\varepsilon > 0$ arbitrarily, which is positive but *close to zero*. Assume that there exist $\Omega_i^* \subset \Omega_i$ and $\Omega_i^{**} \subset \Omega_i$ such that $\Omega_i^* \bigcap \Omega_i^{**} = \phi$,

$$\sum_{\omega_i \in \Omega_i^*} f_i(\omega_i \mid a_i^*) \geq 1 - \varepsilon, \quad \sum_{\omega_i \in \Omega_i^{**}} f_i(\omega_i \mid a_i^{**}) \geq 1 - \varepsilon,$$

$$\sum_{\omega_i \in \Omega_i^*} f_i(\omega_i \mid a_i) \leq \varepsilon \quad \text{for all} \quad a_i \in A_i \setminus \{a_i^*\}, \text{ and}$$

$$\sum_{\omega_i \in \Omega_i^{**}} f_i(\omega_i \mid a_i) \leq \varepsilon \quad \text{for all} \quad a_i \in A_i \setminus \{a_i^{**}\}.$$

This assumption along with a small $\varepsilon > 0$ implies that the signals are *accurate* in monitoring. When player i chooses a_i^*, it is almost certain that the signal ω_i for player i's action choice belongs to Ω_i^*. When player i chooses a_i^{**}, it is almost certain that it belongs to Ω_i^{**}. Hence, opponent j can *almost* perfectly monitor whether player i has chosen a_i^*, a_i^{**}, or other actions.

We specify the function μ_i as

$$\mu_i(\omega_i) = \delta_i\{w_i(a_j^*) - w_i(a_j^{**})\} \quad \text{for all} \quad \omega_i \notin \Omega_i^* \bigcup \Omega_i^{**}, \text{ and}$$

$$\mu_i(\omega_i) = 0 \quad \text{for all} \quad \omega_i \in \Omega_i^*.$$

By selecting $\mu_i(\omega_i)$ from the interval $[0, \delta_i\{w_i(a_j^*) - w_i(a_j^{**})\}]$ for each $\omega_i \in \Omega_i^{**}$ in the appropriate manner, we can make μ_i satisfy (5), (6), and (7), and we can make $\sum_{\omega_i \in \Omega_i} \mu_i(\omega_i) f_i(\omega_i \mid a_i^*)$ close to zero. This implies that the least upper bound $R_i(\delta_i)$ is approximated by $u_i(a^*)$. Hence, we have proved that whenever the signals are sufficiently accurate, efficiency is approximated by a tit-for-tat Nash equilibrium payoff vector even if the discount factors are fixed and not very high.

Theorem 3: *If* $\varepsilon > 0$ *is sufficiently close to zero and (8) holds for each* $i \in \{1,2\}$*, then there exists a tit-for-tat Nash equilibrium* (q,s) *in* Γ *such that* $v(q,s)$ *is approximated by* $u(a^*)$.

6. Example

Fix $p \in (1/2, 1)$ and $\delta \in (0,1)$ arbitrarily. We investigate an example where $A_i = \{0,1\}$, $\Omega_i = \{0,1\}$, $f_i(1|1) = f_i(0|0) = p$, and $\delta_i = \delta$ for each $i \in \{1,2\}$. The following matrix illustrates the component game, where we assume $Z > Y > 0$. Let $a^* = (1,1)$ and $a^{**} = (0,0)$. All the assumptions necessary for the results of this paper are satisfied, i.e., the payoff functions are additively separable, a^* is efficient, and the inequalities expressed in (1) hold.

	1	0
1	X X	X – Z X + Y
0	X + Y X – Z	X + Y – Z X + Y – Z

From the standard calculation, a tit-for-tat strategy profile (q,s) satisfies (4) if and only if

(9) $$s_i(1) - s_i(0) = \frac{Y}{\delta(2p-1)Z} \quad \text{for all} \quad i \in \{1,2\} .$$

This along with Proposition 1 implies that the inequalities given in (9) are necessary and sufficient for the tit-for-tat strategy profile to be a Nash equilibrium. This implies that there exists a tit-for-tat Nash equilibrium (q,s) if and only if

(10) $$\delta \geq \frac{Y}{(2p-1)Z} .$$

Suppose that (10) holds. The standard calculation along with (2) implies that if (q,s) is a tit-for-tat Nash equilibrium, the payoff induced by (q,s) equals

$$v_i(q,s) = X - \frac{1-p}{2p-1}Y - [(1-\delta)(1-q_j) + \delta\{1-s_j(1)\}]Z .$$

The average payoff of the players equals

$$\frac{v_1(q,s)+v_2(q,s)}{2} = X - \frac{1-p}{2p-1}Y - [(1-\delta)(1-\frac{q_1+q_2}{2}) + \delta\{1-\frac{s_1(1)+s_2(1)}{2}\}]Z .$$

Then, it follows that if (q,s) and (q',s') are tit-for-tat Nash equilibria and

$$q_1 + q_2 = q_1' + q_2' \quad \text{and} \quad s_1(1) + s_2(1) = s_1'(1) + s_2'(1) ,$$

then the average payoffs are the same between (q,s) and (q',s'), i.e.,

$$\frac{v_1(q,s)+v_2(q,s)}{2} = \frac{v_1(q',s')+v_2(q',s')}{2} .$$

Since the tit-for-tat Nash equilibrium $(\hat{q},\hat{s})$ specified by (9), $\hat{q}_1 = \hat{q}_2 = 1$, and $\hat{s}_1(1) = \hat{s}_2(1) = 1$ induces the least upper bound $R_i(\delta)$ for each player i, it follows that

$$R_i(\delta) = \tilde{R}_i = v_i(\hat{q},\hat{s}) = X - \frac{1-p}{2p-1}Y .$$

This along with Proposition 2 implies that the least upper bound is independent of the discount factor δ. As Theorem 3 shows, the least upper bound $X - \frac{1-p}{2p-1}Y$ converges to the efficient payoff X as p approaches unity.

In order to demonstrate the degree of implicit collusion *in the long run*, it is appropriate to exclude the payoffs in the early periods and concentrate on the *stationary distribution* defined as $\rho = \rho(s) : A \to [0,1]$, where

$$\rho(a^*) = \sum_{a \in A,\, \omega \in \Omega} \rho(a) s_1(\omega_2) s_2(\omega_1) f(\omega \mid a),$$

$$\rho(a^{**}) = \sum_{a \in A,\, \omega \in \Omega} \rho(a) \{1 - s_1(\omega_2)\}\{1 - s_2(\omega_1)\} f(\omega \mid a),$$

$$\rho(a_i^*, a_j^{**}) = \sum_{a \in A,\, \omega \in \Omega} \rho(a) s_i(\omega_j) \{1 - s_j(\omega_i)\} f(\omega \mid a),$$

and $\rho(a) = 0$ for all $a \notin \{a^*, a^{**}, (a_1^*, a_2^{**}), (a_1^{**}, a_2^*)\}$. Let $\rho_i^* = \sum_{a_j \in A_j} \rho(a_i^*, a_j)$ denote the *relative frequency* of player i's action choice a_i^* in the long run. Note that

$$\rho_i^* = K_j^* \rho_j^* + K_j^{**}\{1 - \rho_j^*\},$$

where $K_j^* \equiv \sum_{\omega_j \in \Omega_j} s_i(\omega_j) f_j(\omega_j \mid a_j^*)$ and $K_j^{**} \equiv \sum_{\omega_j \in \Omega_j} s_i(\omega_j) f_j(\omega_j \mid a_j^{**})$. Hence,

$$\rho_i^* = \frac{K_i^{**}(K_j^* - K_j^{**}) + K_j^{**}}{1 - (K_i^* - K_i^{**})(K_j^* - K_j^{**})}.$$

In this example, the standard calculation along with (9) implies that

$$K_i^* = s_i(1)p + s_i(0)(1-p) = s_i(1) - \frac{Y(1-p)}{\delta Z(2p-1)},$$

$$K_i^{**} = s_i(1)(1-p) + s_i(0)p = s_i(1) - \frac{Yp}{\delta Z(2p-1)},$$

and therefore,

$$\rho_i^* = \frac{\delta Z\{s_j(1)\delta Z + s_i(1)Y\}}{\delta^2 Z^2 - Y^2} - \frac{Yp}{(\delta Z - Y)(2p-1)}.$$

The *average* relative frequency in the long run is expressed as

$$\rho_1^* + \rho_2^* = \frac{\delta Z\{s_1(1) + s_2(1)\}}{2(\delta Z - Y)} - \frac{Yp}{(\delta Z - Y)(2p-1)},$$

which depends only on $s_1(1) + s_2(1)$.

The *least upper bound* of the relative frequencies in the long run is achieved by the tit-for-tat Nash equilibrium $(\hat{q}, \hat{s})$ satisfying $\hat{s}_1(1) = \hat{s}_2(1) = 1$, and is equal to $\frac{\delta Z(2p-1) - YP}{(\delta Z - Y)(2p-1)}$. Since this value is increasing with respect to δ, it follows that implicit collusion in the long run is more successful when the discount factor is higher.

References

Abreu, D., D. Pearce, and E. Stachetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58, 1041–1063.

Ely, J. and J. Välimäki (2002): "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory* 102, 84–105.

Fudenberg, D., D. Levine, and E. Maskin (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica* 62, 997–1040.

Fudenberg, D. and J. Tirole (1991): *Game Theory,* MIT Press.

Green, E. and R. Porter (1984): "Non-cooperative Collusion under Imperfect Price Information," *Econometrica* 52, 87–100.

Osborne, M. and A. Rubinstein (1994): *A Course in Game Theory,* MIT Press.

Piccione, M. (2002): "The Repeated Prisoners' Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory* 102, 70–83.

Sekiguchi, T. (1997): "Efficiency in Repeated Prisoners' Dilemma with Private Monitoring," *Journal of Economic Theory* 76, 345–361.